

Description

[SYSTEM AND METHOD OF ANALYZING TIMING EFFECTS OF SPATIAL DISTRIBUTION IN CIRCUITS]

BACKGROUND OF INVENTION

[0001] Field of the Invention

[0002] The invention generally relates to the timing analysis of circuits, and, more particularly to systems and methods for analyzing the timing effects of spatial distribution in circuits.

[0003] Description of Related Art

[0004] In circuit design, one signal may need to arrive at a particular point in a circuit path before another signal. For example, a signal representing a stable data value may need to arrive at a memory element (e.g., a flip-flop or latch) before the clock signal that stores the value in the memory element. Alternatively, a signal representing a change in a data value may need to arrive at a memory element

after the clock signal that stored the last data value. Two paths along which signals propagate so as to arrive at a particular pair of points (e.g., clock and data pins of a flip-flop) in a defined relationship with one another, as in the examples above, are often referred to as racing paths. Each set of racing paths typically includes an early path and a late path. The comparison of a pair of early and late mode signals to determine whether a particular requirement on their relative arrival times is met is called a timing test. An early mode signal or arrival time is the earliest time at which the value on a net or at a point can change from its previous cycle stable value, and a late mode signal or arrival time is the latest time at which the value on a net or at a point can settle to its final stable value for the current cycle.

[0005] Static timing analysis (STA) is a tool used for verification of circuit design and analysis of circuit performance. STA uses delay models to evaluate the delay in circuit paths. Most delay models used for STA employ parameter distributions, for example, parameter distributions for gate length, to define circuit performance. Each parameter distribution is a statistical measure that defines the possible values for each parameter and the relative likelihood of

obtaining those values.

[0006] Three types of distributions are typically defined. The first type of distribution describes the variance of the chip means and defines possible mean values for a parameter that may be applied systematically across the chip. The second type of distribution describes a local mean that varies across the chip and defines the possible variation of the parameter across the chip. The third distribution is centered around the "local" mean and describes all the random, non-chip-dependent and non-position-dependent variations of the parameter.

[0007] In most cases, the three factors defined by the parameter distributions (mean or average value of a parameter, local variation of a parameter, and random variation of a parameter) are considered together, and an engineer designing a circuit considers the best-case and worst-case variation in a parameter. Typically, this would be done by performing an analysis in which a single value for each parameter is chosen from the chip mean distribution (e.g., the fastest for a "fast chip analysis" or the slowest for a "slow chip analysis"). Within this analysis early mode timing values are determined using the "best-case"(fastest) values within the local mean and random component dis-

tributions, and late mode timing values are determined using the "worst-case"(slowest) values within the local mean and random component distributions. Because the engineer considers the worst-case variation in the local mean of a parameter, pessimistic estimates of the delay in a circuit path are typically used (i.e., an unrealistically large difference between early and late mode timing values is used). Such pessimistic evaluations of the delay in circuit paths may lead to overdesigned circuits.

SUMMARY OF INVENTION

[0008] One aspect of the invention relates to a method for analyzing the timing of a circuit. The method comprises determining at least one location information for one or more inputs to a timing test and computing a timing slack for the timing test using the at least one location information.

[0009] Another aspect of the invention relates to a computer-readable medium containing instructions. When executed, the instructions cause the computer to determine at least one location information for one or more inputs to a timing test and a timing slack for the timing test using the at least one location information.

[0010] A further aspect of the invention relates to a method of

analyzing the timing of an integrated circuit. The method comprises identifying an early path and a late path in the integrated circuit. The method also comprises determining a timing slack variation in the early path using location information on one or more elements in the early path and determining a timing slack variation in the late path using location information on one or more elements in the late path, and computing a new timing slack for the early path and the late path by using the timing slack variation in the early path and the timing slack variation in the late path.

BRIEF DESCRIPTION OF DRAWINGS

- [0011] The invention will be described with respect to the following drawing figures, in which like numerals represent like elements throughout the views, and in which:
- [0012] FIGURE 1 is a schematic flow diagram of a method according to embodiments of the invention;
- [0013] FIGURE 2 is a schematic flow diagram of another method according to embodiments of the invention;
- [0014] FIGURE 3A is a schematic diagram illustrating two paths originating at separate nodes and merging at a node;
- [0015] FIGURE 3B is a schematic diagram illustrating the concept of varying the dimensions of a bounding region, as ap-

plied to the paths illustrated in FIGURE 3A;

[0016] FIGURE 4 is a schematic flow diagram illustrating a method for performing certain tasks illustrated in the flow diagram of FIGURE 1; and

[0017] FIGURE 5 is a schematic diagram illustrating two paths and a method of considering location information of elements in the paths in the aggregate according to embodiments of the invention.

DETAILED DESCRIPTION

[0018] In general, the invention provides systems and methods for more accurately predicting the delay in circuit paths by considering the portion of delay that is influenced by the proximity of circuit elements in a path or paths separately from the full delay distribution. Methods according to the invention may be embodied in systems and software programs for delay analysis, and may be used to improve the design of circuits and circuit paths or to identify the cause of timing failures in existing circuits.

[0019] FIGURE 1 is a schematic flow diagram of a method 10 according to an embodiment of the invention for analyzing the timing of a circuit. The circuit analyzed in method 10, and in other methods according to the invention, may be a circuit or integrated circuit of any type or function and,

in particular, may be an existing integrated circuit or a newly-designed integrated circuit.

[0020] Method 10 begins at S12, and continues at S14. At S14, the user defines the cells or elements that are of interest for a particular timing test or evaluation. These cells or elements may be considered in terms of paths (i.e., the chain of gates through which the earliest or latest signal can pass) or logic cones (i.e., the entire set of gates which converge on a timing test). Cells and elements of interest may include, for example, the metal wires of connection gates, the transistors of the gates, or any other structures affecting the generation and propagation of signals in the circuit. In general, cells or elements of interest may be considered, individually or collectively, to be inputs to a timing test. Once the cells or elements of interest are defined, method 10 continues with S16.

[0021] At S16, the user defines location information for the cells or elements of interest that were defined in S14. Location information may refer to the physical location of a particular element in a circuit layout relative to another, in which case the location information may be expressed in terms of a conventional coordinate system, such as the Cartesian coordinate system. However, the term "loca-

tion," as used herein, should not be construed as being limited to indicia of physical location in a circuit layout. In addition to indicia of physical location, abstract spaces that describe location may be defined by considering factors other than physical location.

[0022] For example, an abstract location space and, therefore, the distance between two elements, may be defined by the similarity between the two elements. As one example, gates of the same or similar design would typically have closely correlated delay distributions, and might be considered close together based on the difference between their delay distributions. However, gates of different design (e.g., pass gate circuits and complementary CMOS circuits) may have little delay correlation, and may be considered far apart based on the difference between their delay distributions. In embodiments of the invention, both abstract location and physical location may be taken into account by defining an appropriate abstract space that includes information on physical location. Once location information is defined in S16, method 10 continues and the timing slack for the paths or elements defined in S12 is calculated at S18, taking into account the location information defined in S16. Timing slack may be calcu-

lated as the difference in arrival times between an early and a late path in a particular test, although the particular calculations to determine timing slack for a particular circuit or test may vary and can be determined by those skilled in the art. Timing slack is typically defined so that a negative slack indicates a violation of a timing constraint, and a positive value indicates the margin by which a timing constraint is met. In typical methods according to embodiments of the invention, the location information defined in S16 would be used to calculate a timing slack variation based on the location information, and the timing slack variation would be added to the typical, unadjusted timing slack for the racing paths or logic cone to determine a new, adjusted timing slack which is less pessimistic.

[0023] Once the timing slack has been determined in S18, the method 10 returns at S20. However, method 10 may be implemented in software or other tools for STA, in which case it may be used iteratively, or other tasks related to method 10 may be performed. For example, method 10 could be embodied, at least in part, in an STA analysis program or utility such as EinsTimer (IBM Corporation, Armonk, New York), which may be used in the course of

logic editing, logic synthesis and placement tools. If used as a part of placement tools, the results of method 10 may be used to apply design penalties for widely located circuit elements according to layout rules, so as to encourage closely located circuit elements, particularly for paths or logic cones whose delay is important to the function of the circuit as a whole.

[0024] As will be understood by those of skill in the art, method 10, and other methods according to the invention, may be implemented on a general purpose computer, a special purpose computer, an ASIC, or any other device having sufficient capabilities to execute the tasks of the methods. Additionally, method 10 and other methods according to the invention may be implemented using a compiled programming language, such as C or C++, an interpreted or scripting language such as Tcl, and other languages, such as Java and J++. Typically, the particular hardware platform and programming language would be chosen so as to complement the hardware and programming languages in which existing STA tools are implemented.

[0025] Each of the three major tasks S14, S16, S18 defined in method 10 may be implemented in a number of ways. For example, FIGURE 2 is a schematic flow diagram of a

method 50 according to an embodiment of the invention. Method 50 begins at S52 and continues with S54. At S54, a list is made of all cells or elements in the path or logic cone of interest and their location coordinates. (As was described above, the location coordinates may be physical, abstract, or some combination of physical and abstract.) Method 50 then continues with S56, in which the list of S54 is traversed to find the maximum and minimum location coordinates in the list. For example, if Cartesian coordinates are used for location, S56 would involve finding the minimum and maximum X and Y coordinates. Depending on the particular circuit path, there may be any number of maxima and minima.

[0026] Once the maxima and minima are determined, method 50 continues with S58, in which a bounding polygon or region is formed using each set of maxima/minima coordinates as an endpoint. The region may have any number of sides and any number of endpoints. Additionally, it should be understood that while the term "polygon" may be used for convenience in description, the sides or distances between maxima and minima may be of any length and, particularly, may be of irregular lengths.

[0027] In some embodiments, particularly those in which method

50 is implemented as part of a software tool with a graphical user interface (GUI), S54–S58 may be performed by, for example, selecting individual elements or groups of elements to define the bounding region or polygon using the features of the GUI. Alternatively, if a GUI is used, the user could select a larger area and the software package could follow S54–S58 to define a bounding region within that larger area.

[0028] The bounding region defined in S58 may be used to calculate delays and adjust timing slack without any further tasks and, in that case, method 50 may continue directly with S62. In S62, the delay due to the relative location of elements within a circuit path may be defined using the bounding region or polygon by defining a function that expresses the relationship between the location of two elements and their delay difference. This task will be described in more detail below.

[0029] However, if the bounding region defined in S58 is used, the location-based adjustments to the timing slack may be overly pessimistic, because the bounding region may encompass the entire logic cone. For that reason, in embodiments of the invention, method 50 continues with S60, in which the bounding region or polygon is modified

to make a more realistic estimate of the location-based adjustments to the timing slack.

[0030] FIGURES 3A and 3B are schematic diagrams illustrating the basic concept of modifying the bounding region. FIGURE 3A illustrates a path A, which originates from a node or cell 100 and terminates/merges at a node or cell 102, and a path B, which originates from a node 104 and terminates/merges with path A at node or cell 102. As an example, path A may have an arrival time of 10 and path B may have an arrival time of 5. (For purposes of this conceptual example, units of distance and arrival time are arbitrary.) FIGURE 3B illustrates an appropriate bounding region for path A, indicated by reference numeral 106, and an appropriate bounding region for path B, indicated by reference numeral 108. The two bounding regions 106, 108 encompass the elements or cells in the paths A, B.

[0031] Assume that a distance of X between two objects of interest can introduce at most an increased delay difference of D. In the example of FIGURE 3A, the arrival time propagated along A is the extreme (e.g., the maximum for late mode), and the arrival time propagated along B is D less than that propagated along A (in the case of FIGURE 3A, $D=5$). Because of this, the bounding shape 108 of path B

may be reduced in each dimension by X , as shown in FIGURE 3B, to form a reduced B shape 110 before merging the reduced B shape 110 with the bounding shape 106 for path A to form a merged bounding shape 112.

[0032] The distance from any point in this merged bounding shape 112 to another point in a bounding shape for the opposite path (e.g., the early path) will still determine a conservative bound on the slack at the test. This can be seen by considering two cases. First, the farthest point in the merged A/B shape 112 may come from the A shape 106. Because the A shape 106 bounds the locations of objects of interest feeding A, a delay variation DV computed from it will also bound the delay variation of the difference between the early path arrival time and the late path arrival along A. Second, the farthest point in the merged A/B shape 112 may come from the reduced B shape 110. Because the selected point is the furthest point from the early path of any in the merged A/B shape 112, no object of interest in the B path may be more than X farther from the early path than this selected point. Therefore, the delay variation between an object in path B and one in the early mode path will be at most $DV + D$. However, since the nominal arrival time along A is already

D greater than that along B, adding D to the variation of B cannot create a worse slack at the test.

[0033] The reduction of S60 that was explained above may be performed on both early and late paths at every point where multiple delay edges converge, where the shape for the non-dominant edges (the edges which do not produce the maximum arrival time in late mode, or the minimum arrival time in early mode) are reduced. The reduction distance (the shrinkage factor) is determined by the cumulative delay along the path to the point, and the difference in arrival time from the dominant arrival time. The cumulative delay influences the reduction distance because delay variation is a fraction of total delay, and so the reduction distance for a sub-cone with a large cumulative delay will be smaller for a given delta arrival time than for a sub-cone with a small cumulative delay.

[0034] The tasks of S60 may be represented by a set of mathematical equations and relationships, which may facilitate computerized or other automated implementation of methods according to the invention. With respect to FIGURE 3A, the two nodes of interest, 100 and 104, have outgoing timing edges A and B, respectively, both of which terminate on 102, and the late mode bounding re-

gions and arrival times for 100 and 104 (without considering variation) may be defined as $L_{box}(100)$, $L_{box}(104)$, $LAT(100)$, and $LAT(104)$ (with "AT" denoting "arrival time"). (Of course, the equations and relationships may be defined with respect to any nodes of interest; reference numerals from FIGURE 3A are used only for convenience in description.) The maximum late mode delays of edges A and B (without considering variation) may be defined as $L_d(A)$ and $L_d(B)$. By STA:

[0035]
$$LAT(102) = \max(LAT(100) + L_d(A), LAT(104) + L_d(B)) \quad (1)$$

[0036] Or, in other words, the arrival time of 102 is the maximum of either the arrival time of 100 plus the maximum late mode delay of edge A or the arrival time of 104 plus the maximum late mode delay of edge B. For purposes of this example, assume that the $LAT(100) + L_d(A)$ term of Equation (1) is the larger of the two terms. With that assumption, a function L_{diff} may be defined as:

[0037]
$$L_{diff} = LAT(100) + L_d(A) - (LAT(104) + L_d(B)) \quad (2)$$

[0038] A function may also be defined to relate distance to delay. For example, a function may be defined such that elements at the same location would have a 0% delay difference, elements separated by 1 mm or less would have a

3% delay difference, elements separated by 2 mm or less would have a 5% delay difference, and so forth. The data used to define the distance–delay function may be gathered empirically or estimated, and the type of distance–delay function may vary among embodiments of the invention. For purposes of this description, the distance–delay function described above may be called $\text{var}(d)$. Additionally, a function $\text{var_inv}(p)$ is defined as the inverse of the function $\text{var}(d)$ (i.e., $\text{var_inv}(p)$ maps a percentage delay difference to a distance), a variation in the bounding box or region may be calculated.

[0039] Because, in this example, the $\text{LAT}(100) + \text{Ld}(A)$ term of Equation (1) is the greater of the two terms, $\text{Lbox}(104)$ can be reduced. To calculate the area reduction distance for $\text{Lbox}(104)$, Equation (3) may be used:

[0040] $\text{reduction}(104) = \text{var_inv}(p)$ (3)

[0041] In which the variable p is defined by Equation (4):

[0042] $p = \text{Ldiff} / (\text{LAT}(104) + \text{Ld}(B))$ (4)

[0043] Using Equations (3) and (4), $\text{Lbox}(104)$ is shrunk by $\text{reduction}(104)$ by adding $\text{reduction}(104)$ to the left edge X coordinate and bottom edge Y coordinate, and subtracting $\text{reduction}(104)$ from the right edge X coordinate and the

top edge Y coordinate. Finally, to generate a new bounding region, the union of Lbox(100), Lbox_reduced(104), and the region enclosing 102 is taken. (Optionally, the output net of 102 may also be included in the union.) To do this, a new bounding region is computed in which, for example, the left edge X coordinate and bottom edge Y coordinate are the minima of those of the regions included in the union and the right edge X coordinate and bottom edge Y coordinate are the maxima of those of the regions included in the union.

[0044] With the calculations described above, it is possible that an area reduction may cause Lbox_reduced(104), or another bounding region for a non-dominant path, to have a negative extent (e.g., have a left edge X coordinate greater than its right edge X coordinate, or a bottom edge Y coordinate greater than its top edge Y coordinate). This is acceptable, and the union calculations described above can still be performed. Because neither the bounding box the cell on which the paths are converging (in this case 102) nor that for the dominant path (in this case the path from 100) can have negative extent, it follows that the extent of Lbox(102) generated via the union operation cannot be negative.

[0045] The calculations are essentially the same for the early mode analysis. For early mode analysis, the early mode bounding regions and arrival times for 100 and 104 (without considering variation) may be defined as $E_{box}(100)$, $E_{box}(104)$, $EAT(100)$, and $EAT(104)$, and the minimum (early mode) delays of edges A and B (without considering variation) be $Ed(A)$ and $Ed(B)$. By STA:

[0046] $EAT(102) = \min(EAT(100) + Ed(A), EAT(104) + Ed(B))$ (5)

[0047] Without loss of generality, assume that

[0048] $EAT(100) + Ed(A) > EAT(104) + Ed(B)$ (6)

[0049] and that

[0050] $Ediff = EAT(100) + Ed(A) - (EAT(104) + Ed(B))$ (7)

[0051] The reduction process is then the same for early mode analysis as it is for late mode analysis, and $E_{Box}(100)$, being the noncritical or non-dominant path, is reduced by $var_inv(Ediff / (EAT(100) + Ed(A)))$.

[0052] Once the bounding region is modified for variability in S60, method 50 continues with S62, in which a distance-delay function is defined to convert from units of distance to units of delay. For example, the functions $var(d)$ and $var_inv(p)$, which were defined above, may be used to re-

late distance to delay once the appropriate bounding regions are calculated, so that a timing slack variation can be calculated. The timing slack variation would typically be added to the original timing slack for the paths. Once an improved timing slack is calculated, method 50 terminates and returns.

[0053] Method 50 provides one way to evaluate the location-based delays in a circuit. However, methods according to the invention need not perform the precise tasks of method 50. For example, in S54 and S56 of method 50, elements of interest may be defined by traversing a list that includes the location coordinates of each element. Alternatively, elements of interest could be defined, for example, by a method of beginning at one element or node in a circuit and traversing the paths leading to and emanating from that element or node.

[0054] FIGURE 4 is a schematic flow diagram of an alternative method for defining elements of interest (i.e., S14 of method 10), shown with respect to the tasks of method 10. However, it should also be understood that the tasks illustrated in FIGURE 4 may replace S54–S58 of method 50. The tasks illustrated in FIGURE 4 begin when control is transferred from S12 of method 10.

[0055] The first task illustrated in FIGURE 4 is S140, in which either a user or an automated algorithm begins at a particular node and traverses a path leading to or emanating from that element or node. Although the remaining tasks illustrated in FIGURE 4 will be described with respect to a single path, multiple paths may be traced simultaneously or in parallel, limited only by the ability of the user or algorithm to do so. As each element in the path is encountered, control passes to a decision task, S142, in which the user or algorithm decides whether an element encountered in the path is of interest for timing analysis purposes. An element would be of interest, for example, if it is part of the path or paths that are to be analyzed. An element may also be of interest if it feeds or leads into a path that is part of a path or paths that are to be analyzed. However, in methods according to the invention, limits may be set on how many secondary or feeding paths are included in the bounding region. An element would typically be considered to be of interest if it is on at least one path which fails to meet a timing constraint (e.g., has a negative slack) according to a conventional static timing analysis. For other elements, the pessimism involved in a conventional static timing analysis may be

less important, since the paths on which they appear will meet their timing constraints even without elimination of this pessimism. An element may be also of interest if its delay is highly dependent on a parameter which is known to have significant spatial correlation. Such high dependence might be identified, for example, by determining whether the magnitude of the derivative of the delay of the element with respect to a parameter value exceeds some threshold.

[0056] If the determination is made in S142 that an element is of interest (S142:YES), control passes to S146, in which the element is added to a list of elements that are of interest, before passing to S148, in which the location of the element is noted. The location noted in S148 may be the physical location of the element in the circuit (e.g., in X and Y coordinates), looked up, for example, in a database containing location information for each element or calculated from the starting node plus the path distance traversed. Alternatively, as described above, the location noted in S148 may be an abstract location calculated by examining one or more of characteristics of the element, such as the type of circuit (e.g., pass gate, static CMOS, dynamic), characteristics of the element (e.g., the stack

height, of number of series transistors in the function), the delay function of the element (e.g., the parameters on which the delay depends and the sensitivity of the delay to each parameter), or other considerations, either with or without also considering the physical location. The path tracing involved in selecting the elements of interest may also be performed separately from the actual selection, as part of a conventional static timing analysis in which slack values are determined for each element in an integrated circuit. Selection of elements of interest in this case may simply comprise selecting elements whose conventionally computed slack is less than some threshold.

[0057] If the determination is made in S142 that an element encountered in the path is not of interest (S142:NO), control passes to S144, a decision task, in which the user or algorithm decides whether the end of the path has been reached. If the end of the path has been reached (S144:YES), control passes out of the tasks illustrated in FIGURE 4, for example, into S16 of method 10 or S60 of method 50, depending on the particular method that is implemented. If the end of the path has not been reached (S144:NO), control returns to S140 and the user or algorithm continues to traverse the path. S150, to which con-

trol passes from S148, is a decision task identical to S144, and branches in the same ways as S144.

[0058] In the methods and description above, location information was considered for each individual element or cell in a path or logic cone. However, in embodiments of the invention, location of elements or cells in a path or logic cone may be considered in the aggregate, instead of individually.

[0059] FIGURE 5 is a schematic diagram illustrating a method of considering location information in the aggregate. In FIGURE 5, an early path 200 and a late path 202 are shown schematically, laid out on a Cartesian coordinate plane with an origin at (0,0). The early and late paths 200, 202, merge at a latch 204 or other element. The early path 200 includes elements or cells 206, 208, 210, and 214; the late path 202 includes elements 216, 218, and 220. Each element 206–220 has its own unique location in the circuit, as shown in FIGURE 5, and each element 206–220 has some characteristic delay. In this example, elements 206, 208, and 220 have delays of 1 ns, element 214 has a delay of 4 ns, and elements 216 and 218 have delays of 2 ns, 5 ns, respectively. The delays shown in FIGURE 5 are for illustrative purposes; in general, delays considered

may include individual cell delays alone, or cell delay plus the wire delay of the cell's output net.

[0060] Based on the locations of each cell or element in the paths 200, 202 and, optionally, their delays, a centroid or "delay center of gravity" is calculated for each one of the paths individually. FIGURE 5 shows the two centroids, including the centroid for the early path C_E and the centroid for the late path C_L . The distance D between the two centroids C_E , C_L and, optionally, the standard deviations of the centroid locations, may be used to calculate a timing slack variation, for example, by using the type of distance-delay functions described above with respect to method 50.

[0061] The centroids may be calculated in a variety of ways. In one method of calculating a centroid, the location coordinates of each element in a path would simply be averaged. However, as was described above, the cell delay or some combination of cell and wire delay may also be taken into account for each cell or element. For example, if the X coordinate, Y coordinate, and delay of a cell or element were to be represented by the variables x_i , y_i , and d_i , respectively, then the X and Y coordinates of the delay-weighted centroid, denoted by C_x and C_y , could be calculated as in Equations (8) and (9):

$$C_x = \frac{\sum_i (x_i \times d_i)}{\sum_i d_i} \quad (8)$$

$$C_y = \frac{\sum_i (y_i \times d_i)}{\sum_i d_i} \quad (9)$$

[0062] As those of skill in the art will realize, a cell or element consumes some non-zero amount of space itself. Therefore, in methods according to the invention, the location coordinates of a cell or element may be the coordinates for the center of the cell. Since most cells or elements are rectangular, this location is unambiguous for most cells. For non-rectangular cells, the cell center could be considered to be the center of the minimum size axis-aligned rectangular bounding box which completely encloses the cell. In general, the center of the cells or elements in a path or logic cone may be defined in any desired way, as long as usage is consistent among multiple cells or elements, or any inconsistencies are taken into account.

[0063] The centroid calculations described above and with respect to FIGURE 5 are only one way in which aggregate and average location information may be used in methods according to embodiments of the invention. Other methods of using aggregate and average location information

may be used in methods according to embodiments of the invention.

[0064] While the invention has been described with respect to certain exemplary embodiments, those of skill in the art will realize that modifications and variations are possible within the scope of the appended claims.